

REMARKS

This Preliminary Amendment is filed in response to the FINAL Office Action mailed on December 4, 2001, and is also filed in the Continuing Prosecution Application filed on even date herewith. All objections and rejections are respectfully traversed.

Claims 1-3, 9, 11, 12, 15-18, 20, 21, 23, 26, 27, 29-32, 44-96 are in the case.

Claims 45- 96 were added to better claim the invention.

Claim 2 is pending, although it is not mentioned in the Office Action dated December 4, 2001 at line 4.

At paragraph 1 of the Office Action it was indicated that documents mentioned in an Information Disclosure Statements (IDS) filed on May 10, 1999 and June 16, 1999 had not been submitted to the United States Patent and Trademark Office.

Applicant notes that the required documents were filed in an IDS on January 10, 2002. Also, a second IDS submitting different documents was filed on January 15, 2002.

At paragraph 2 of the Office Action Claim 25 was objected to. Claim 25 has been cancelled without prejudice.

At paragraphs 3 and 4 claims 1, 11, 25 were rejected under 35 U.S.C. 112, first paragraph. Amendment of the claims is believed to obviate this rejection.

At paragraph 5 claims 15 and 16 were rejected under 35 U.S.C. 112, second paragraph. Amendment of the claims is believed to obviate this rejection.

At paragraph 6 of the Office Action claims 1, 2, 17, 18, 20, 21, 23, 24-27, 29-35, and 37-44 were rejected under 35 U.S.C. 102(e) in view of Imai et al. U. S. Patent No. 6,175,874 (hereinafter Imai patent).

The present invention, as set forth in representative claim 1, comprises in part:

1. A routing system for distributing packets in a network, wherein the packets originate at a source and are routed to a destination, comprising:
 - a plurality of route processing engines;
 - a mechanism that performs a hashing function on at least a portion of network layer information in the packets transferred to the routing system, to produce an indicia of a flow and,*
 - means for switching packets with a same said indicia of a flow to a single route processing engine.*

The present invention, as set forth in representative claim 26 comprises, in part:

26. A system, in a router, for selecting one processing engine of a plurality of processing engines for processing at least one packet, the system comprising:
 - means for examining at least a portion of network layer flow information of the at least one packet; and*
 - means, responsive to said at least a portion of network layer flow information, for selecting the one processing engine to preserve packet flow indicated by the at least a portion of network layer flow information.*

The present invention, as set forth in representative claim 45, comprises in part:

45. A router, comprising:
 - a plurality of processing engines for processing packets;
 - an interface for receiving a received packet from a network;

a data compiler to perform a hash function on said received packet to generate a hash result, and to select a selected processing engine from said plurality of processing engines in response to said hash result; and,
a switch to distribute said packet to said selected processing engine.

The Imai patent discloses a computer network having a plurality of processing nodes, and having a relay device, where the relay device has a network address corresponding to a node which processes requests, referred to as a virtual node. The relay device distributes requests which it receives from the computer network to a plurality of processing nodes, and each of the processing nodes performs a response to the request addressed to the virtual node. The relay device computes a hash function based on data in a received request packet to generate a hash output, and uses the hash output to select a processing node, selected from the plurality of processing nodes, to process the request. The hash function is used instead of interrogating a Domain Name Service (DMS) located elsewhere in the computer network to find an address of one of the plurality of processing nodes, partly to overcome the DMS problem of having multiple processing nodes offering the same service. The processing nodes form a cluster of computers where each performs the same service, and the relay device maintains an IP address (virtual IP address) so that a user can transparently use one of the nodes of the cluster without knowing or caring which computer of the cluster he used (Col. 4, lines 35-41).

Applicant respectfully urges that applicant's **router** having;

a mechanism that performs a hashing function on at least a portion of network layer information in the packets transferred to the routing system, to produce an indicia of a flow and,
means for switching packets with a same said indicia of a flow to a single route processing engine

as set out in representative Claim 1, and having;

a data compiler to perform a hash function on said received packet to generate a hash result, and to select a selected processing en-

*gine from said plurality of processing engines in response to said hash result; and,
a switch to distribute said packet to said selected processing engine*

as set out in representative Claim 45 is nowhere disclosed by the Imai patent.

Applicant claims a router having a plurality of processing engines to perform a variety of tasks (as set out in various dependent claims) on packets being routed by the router, and a “data compiler” section to select the processing engine by use of a hash function performed on part of the packet received for routing.

Applicant’s router is a single box device which performs routing on packets addressed to a destination node elsewhere in the computer network. Applicant respectfully urges that the complex network cluster described by Imai cannot anticipate Applicant’s claimed single box device under 35 U.S.C. 102.

At paragraphs 7 and 8 of the Office Action claims 9 and 11 were rejected under 35 U.S.C. 103(a) as being unpatentable over Imai.

Claim 9 is dependent upon claim 1, claim 1 is believed to be in condition for allowance, and accordingly claim 9 is believed to be in condition for allowance.

The invention, as set forth in representative claim 11, comprises in part:

11. A router for distributing packets in a network, wherein the packets originate at a source and are routed to a destination, comprising:
 - a plurality of network interfaces that transfer the packets to a destination and from a source;
 - a plurality of route processing engines;
 - a fabric interconnecting said plurality of network interfaces and said plurality of route processing engines;

a hashing function to determine a distribution of the packets, by said fabric in response to an output of said hashing function, among said plurality of route processing engines.

Applicant respectfully urges that Applicant's claimed router having:
a hashing function to determine a distribution of the packets, by said fabric in response to an output of said hashing function, among said plurality of route processing engines

as set out in representative claim 11 is totally absent from the Imai patent.

Further, Applicant respectfully urges that the Imai patent is non-analogous art, as Imai refers to switching packets to different processing nodes of a cluster of processing nodes connected to a computer network. In sharp distinction, Applicant claims a router, with switching inside the router. The router is a single box device which routes packets through the computer network. Inside the router, different processing engines for routing and processing packets which arrive at the router are selected. A single box router is totally distinct from the computer network having a cluster of nodes connected to a computer network, as disclosed by Imai.

Still further, Applicant points out that the Imai disclosure is for a one-to-many distribution, where packets with the Internet address (IP address) of his relay device are translated into the many Internet addresses of computers in his cluster of many computers.

In sharp contrast, Applicant's claimed novel router performs a many-to-many distribution scheme as packets with many different Internet destination addresses are hashed and distributed within Applicant's router to processing engines within Applicant's single box router. That is, Applicant's hash function has many IP destination addresses as input, and produces hash outputs dependent upon these many incoming IP destination addresses. The hash outputs of Applicant's hash functions then perform a function which

is totally beyond any disclosure of Imai. That is, Applicant's hash outputs depend upon the different IP **destination** addresses, different destination port, same protocol, etc. found in packets received by Applicant's claimed novel router.

The different **destination** addresses, different destination ports, same protocol, etc., in packets received by Applicant's router then affect Applicant's hash results, and help to define a flow to which the received packets belong. And by sending all packets having the same hash result to the same processing engine, Applicant switches all packets having the same flow, as determined by the hash function, to the same processing engine. By switching, in response to a flow determined in part by the different destination addresses, different destination ports, same protocol, etc., in the packets, Applicant avoids the problem set out in the Background (p. 3 lines 9-10, and p. 5 lines 3-4) of the present Patent Application, that of routing packets out of order by the router.

If different packets of a single flow were switched to different processing engines, each processing engine would most probably introduce a different delay because of different filling in its queues, etc. The different delays would put the different packets into a different order as they leave the router, and so the packets would have a different order as they are received by the destination computer indicated by the destination IP address of the packet.

For example, a simple round robin distribution of packets to different processing engines within the router would put packets of a single flow into different processing engines. Distributing packets of a single flow into different processing engines would very likely lead to reception of the packets out of order by the computer with the destination IP address. This out of order problem is solved by Applicant's claimed novel use of a hashing function to sort different packets of the same flow into the same processing engine. And, the Imai disclosure has no ability to sort packets of different destination addresses, destination ports, same protocol, etc., into the same processing engine as does

Applicant's claimed novel invention because all of Imai's packets have the same destination IP address.

Accordingly, Applicant respectfully urges that the Imai patent is legally precluded from anticipating Claim 11, or any other claim, under 35 U.S.C. 103(a) because of the absence of any disclosure of a router having the claimed hashing function to determine the distribution of packets by a fabric to a processing engine located inside the router.

At paragraph 9 of the Office Action claims 3 and 12 were rejected under 35 U.S.C. 103(a) as being unpatentable over Imai in view of Varghese et al. U. S. Patent No. 5,905,723.

The Varghese patent discloses a router having a multiport switch connected to ports through links. The links are divided into hunt groups. In the event that a packet is to be routed to a port of the same hunt group through which the packet arrived at the router, then the multiport switch sends the routed packet back through the port from which it arrived, thereby avoiding the packet passing through the backplane of the router.

Applicant respectfully urges that the Varghese patent does not disclose Applicant's claimed novel **router** having:

*a mechanism that performs a hashing function on at least a portion of network layer information in the packets transferred to the routing system, to produce an indicia of a flow and,
means for switching packets with a same said indicia of a flow to a single route processing engine*

as set out in representative claim 1.

Particularly, Applicant respectfully urges that neither Imai nor Varghese disclose a router using a hash function to select a processing engine from a plurality of processing engines located inside of the router.

Accordingly, Applicant respectfully urges that the Imai patent and the Varghese patent, taken either singly or in combination, are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. 103, because of the absence in both the Imai and the Varghese patent of Applicant's claimed novel use in a router of a hash function to select a processing engine from a plurality of processing engines.

Further, Applicant respectfully urges that Applicant's representative claim 26 is not anticipated under 35 U.S.C. 102 nor made obvious under 35 U.S.C. 103 by any cited art, either taken singly or in any combination, because of the absence in any of the cited art of Applicant's claimed novel router having:

means for examining at least a portion of network layer flow information of the at least one packet; and
means, responsive to said at least a portion of network layer flow information, for selecting the one processing engine to preserve packet flow indicated by the at least a portion of network layer flow information.

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account
No. 03-1237.

Respectfully submitted,

A handwritten signature in cursive script, appearing to read "A. Sidney Johnston", written over a horizontal line.

A. Sidney Johnston
Reg. No. 29,548
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500

**MARK-UP PAGES FOR THE APRIL 1, 2002, AMENDMENT TO
U.S. PATENT APPLICATION SER. NO. 09/053,237**

The replacement for claims resulted from the following changes:

1. (Amended) 1. (four times amended) A routing system for distributing packets in a network, wherein the packets originate at a source and are routed [returned] to a destination,

[both source and destination external with respect to the routing system,]

comprising:

[a plurality of means for transferring packets to a destination and from a source;]
a plurality of route processing engines; [and]

a mechanism that performs a hashing function on at least a portion of network layer information [,] in the packets transferred to the routing system, to produce an indicia of a flow [determine an approximately even distribution of the packets to the route processing engines for processing by the engines];

means for determining packets belonging to a same flow;

means for determining the original order of the packets in the same flow from the network layer information of the packets, the network layer information including at least the same source/destination and protocol] ; and,

means for switching packets with a same said indicia of a flow [preserving the original ordered packet flows by sending each ordered packet flow] to a single route processing engine.

2. (thrice amended) The routing system of claim 1,

[wherein the plurality of means for transferring packets includes]

further comprising:

at least one fast uplink connection to an external network to accept outgoing packets from a plurality of processing engines

[and at least one data port adapter connected to an external data interface component, wherein the port adapter converts s input data to a known interface, and the external interface component accepts input data].

3. (Amended) The [scalable] routing system of claim 1, further comprising:
[also including]

a crossbar as said means for switching packets.

9. (Amended) The routing system of claim 1, further comprising:
[wherein]

means for scaling processing power of said system [can be scaled] by adding additional route processing engines to said plurality of route processing engines.

11. (thrice amended) A router [routing system] for distributing packets in a network, wherein the packets originate at a source and are routed
[returned]

to a destination,
[both source and destination external with respect to the routing system,]
comprising:

a plurality of network interfaces that transfer the packets to a destination and from a source;

a plurality of route processing engines;

a fabric interconnecting said plurality of network interfaces and said plurality of route processing engines;

[wherein each of said plurality of network interfaces uses]

a hashing function to determine a distribution of the packets, by said fabric in response to an output of said hashing function, among said plurality of route processing engines [; and]

[wherein the hashing function is carried out on at least a portion of network layer information in the packets; and

wherein the hashing function determines packets belonging to a same flow;

wherein the hashing function determines the packets' original order from the network layer information including at least the same source/destination and protocol; and means for preserving the original ordered packet flow by sending the original ordered packet flow to a single route processing engine].

12. (Amended) The [scalable] routing system of claim 11, further comprising:

[wherein] said fabric includes a crossbar.

15. (twice amended) The routing system of claim 11, further comprising:

[wherein said network interfaces include]

a port [adapters] adapter, wherein the port adapter converts [s] input data to a desired [known] interface.

16. (twice amended) The routing system of claim 15, wherein said network interfaces include at least one uplink connection to an external network, wherein the port adapter converts [s] input data to a known interface.

17. (twice amended) A method, in a router, for selecting one processing engine of a plurality of processing engines for processing at least one packet, the method comprising the steps of:

hashing at least a portion of network layer information of at least one packet to determine a hash result, said hash result indicating a flow;

[distribution of the packets to be sent to the processing engines;
determining from the network layer information, including at least the source/destination and protocol, the at least one packet that belongs to an ordered packet flow, and]

selecting one processing engine to process the flow indicated by said hash result

[at least one packet that belongs to an ordered packet flow, thereby preserving the ordered packet flow].

18. The method of claim 17, further comprising:

[wherein] the network layer [flow] information comprises one or more of the following network information: a network source address of the at least one packet, a network destination address of the at least one packet, a source port of the at least one packet, and a protocol type value of the at least one packet.

20. (thrice amended) The method of claim 17, further comprising:

[wherein] the hashing is computed by logically XORing an [the] addresses, a [the] port, and a [the] protocol type value.

21. (amended) The method of claim 17, further comprising [the steps of]:

providing a table containing entries for use in selecting the one processing engine;
[and]

selecting [selection] one entry in the table specified by an index value, the index value [being] based upon the hash value to select the processing engine for the hash value

[, and,

using the index value to direct the selection of the one processing engine for those [related] packets that belong to the same packet flow].

23. The method of claim 17, further comprising:

[wherein]

distributing, in response to the hash function, the packets evenly among the plurality of

[the at least one packet is the one of a plurality of packets, and the step of hashing is performed using a hashing function that causes the packets to be at least mostly evenly distributed among the]

processing engines.

26. A system, in a router, for selecting one processing engine of a plurality of processing engines for processing at least one packet, the system comprising:

means for examining at least a portion of network layer flow information of the at least one packet; and

means, responsive to said at least a portion of network layer flow information, for selecting the one processing engine

[based upon, at least in part, the portion of the network layer flow information in such a way as]

to preserve [an original] a packet flow indicated by the at least a portion of network layer flow information

[comprising the at least one packet].

27. The system of claim 26 wherein the network layer flow information comprises:
[one or more of the following network information:]

at least one of

a network source address of the at least one packet,
a network destination address of the at least one packet,
a source port of the at least one packet,
a destination address of the at least one packet, and
a protocol type value of the at least one packet.

29. (amended) The system of claim 26, further comprising:

[wherein]

the hash value is computed by logically XORing the addresses, the ports, and the protocol type value.

30. (amended) The system of claim 26 further comprising:

means for providing a table containing entries for use in selecting the one processing engine; and

means, responsive to the hash value, for selecting one entry in the table
[specified by an index value, the index value being based upon the hash value].

31. The system of claim 26, further comprising:

[wherein the at least one packet is one of a plurality of packets in at least one original flow, and]

the means for selecting [hashing] carries out a hashing function that preserves the packet flow [at least one original flow of the packets].

32. The system of claim 26, further comprising: [wherein]

the at least one packet is one of a plurality of packets, and the means for selecting [hashing] carries out a hashing function that causes the packets to be [least] mostly evenly distributed among the processing engines.

44. A routing system for distributing packets in a network, wherein the packets originate at a source and are routed [returned] to a destination, both source and destination external with respect to the routing system, comprising:

a plurality of network interfaces that transfer packets to said [a] destination and from said [a] source;

a plurality of route processing engines;

a hash mechanism that performs a hashing function on at least a portion of network layer information of a particular packet, in the packets transferred to the routing system, to determine an approximately even distribution of the packets to the route processing engines for processing by the engines, and said hash mechanism producing a hash result giving an indication of a flow of said particular packet so that packets of a flow are switched to the same route processing engine of said plurality of route processing engines

[a processing mechanism that determines packets belonging to a same flow and their original order from the network layer information of the packets, the network layer information including at least the same source/destination and protocol,

a data transferer that sends each ordered packet flow to a single route processing engine, thereby preserving the original ordered packet flows].